arXiv:1612.04988v1 [cs.CL] 15 Dec 2016

TeKnowbase: Towards Construction of a Knowledge-base for Technical Concepts

Prajna Upadhyay
IIT Delhi

Tanuma Patra
IIT Delhi

Ashwini Purkar
IIT Delhi

Maya Ramanath
IIT Delhi

Abstract

In this paper, we describe the construction of TeKnowbase, a knowledge-base of technical concepts in computer science. Our main information sources are technical websites such as Webopedia and Techtarget as well as Wikipedia and online textbooks. We divide the knowledge-base construction problem into two parts – the acquisition of entities and the extraction of relationships among these entities. Our knowledge-base consists of approximately 100,000 triples. We conducted an evaluation on a sample of triples and report an accuracy of a little over 90%. We additionally conducted classification experiments on StackOverflow data with features from TeKnowbase and achieved improved classification accuracy.

1 Introduction

As digital information gains more and more prominence, there is now a trend to organize this information to make it easier to query and to derive insights. One such trend is the creation of large *knowledge-bases* (KBs) – repositories of crisp, precise information, which are machine readable. Indeed the creation of such knowledge-bases has been a goal for decades now with projects such as Cyc [11] and Wordnet [14].

With advances in information extraction research, and the availability of large amounts of structured and unstructured (textual) data, *automatic* construction of knowledge-bases are not only possible, but also desirable because of the coverage they can offer. There are already many such general-purpose knowledge-bases such as Yago [19] and DBPedia [10]. Moreover, projects such as OpenIE [2] and NELL [3] aim to extract information from unstructured textual sources on a large scale.

However, as the technology to automatically build large knowledge-bases matures, there is a paucity of high quality and *specialized* KBs for specific domains. For some domains, for example, for the bio-medical domain, there are well-curated ontologies which partially address this gap (see, for example, the Gene Ontology project [1]). However, for domains such as Computer Science or IT in general, where such curation efforts are hard and the field itself is rapidly growing, it becomes critical to revisit the automatic construction processes that take advantage of domain-specific resources.

In this paper, our aim is to automatically construct a *technical* knowledge-base, called TeKnowBase, of computer science concepts. One of the most important tasks in building any such "vertical" KB is the identification of the right kinds of resources. For example, even though Wikipedia contains technical content, identifying the right subset of this content is crucial. Similarly, while free online technical content, such as textbooks are available, it is important to identify what kind of extractions are possible. The identification of the right resources can sometimes yield bigger gains than using an elaborate information extraction technique.

A preliminary examination of computer science-related resources show that information can be extracted from many different kinds of sources, including, Wikipedia, technical websites such as Webopedia, online textbooks, technical question answer fora such as StackOverflow, etc. By studying

Relation	Example	# Triples
type	Topological_sorting typeOf Graph_algorithms	44,221
concept	⟨Nash_equilibrium conceptsOf Game_theory⟩	833
subTopic	⟨Hamming_code subTopicOf Algebraic_Coding_Theory⟩	1520
application	⟨Group_testing applicationOf Coding_theory⟩	1650
terminology	⟨Blob_detection isTerminology Image_Processing⟩	32,722

Table 1: Statistics for and examples of a subset of relationships

these resources more closely, we developed simple, but effective techniques to build TeKnowbase. Our first step is to acquire a dictionary of concepts and entities relevant to computer science. Using this dictionary, we can further extract relationships among them. We make use of the semantic web standard, RDF, where information is represented as triples of the form ⟨subject⟩⟨predicate⟩⟨object⟩ – in a nutshell, each triple makes a statement about the ⟨subject⟩. Table 1 shows examples of the kind of triples we extract and the number of such triples in our knowledge-base.

Such a knowledge-base, that organizes the space of technical concepts in a systematic way, can be used in a variety of applications. For example, as with general-purpose knowledge-bases, a technical KB can be used for improving classification accuracy, disambiguation of text, linking entities, and, in general, semantic search. Moreover, a technical KB can be of use in a variety of learning scenarios. For example, to students wanting to learn about a new concept, a technical KB can be a valuable resource in identifying related concepts and perhaps pre-requisites. Another useful application is the automatic generation of questions to test general, and conceptual knowledge. For example, a question such as "Name some applications of coding theory" can be generated and the answer graded automatically (see a similar example in [17] in the context of general-purpose KBs).

Note that the set of technical concepts that we consider here are limited to *named* entities. In addition to these, there could be many unnamed entities such as theorem statements, formulae, equations, algorithm listings, etc.[1] In principle, we will be able to recognize these entities and provide a system-generated name to them. But, it is challenging to go beyond this to provide means of querying these entities. In this paper, we limit our scope to only named entities.

Contributions Our contributions are as follows:

- We describe the construction of TeKnowbase, a knowledge-base of technical concepts in computer science. Our approach is general enough that it can be used for other subjects as well.
- In order to construct TeKnowbase, we carefully studied different resources which could be helpful and systematically explored information extraction techniques to collect triples for our KB. We highlight their strengths and drawbacks.
- We present a study on the quality of triples in TeKnowbase that shows a precision of about 90%.
- We perform a simple classification experiment using features from TeKnowbase to improve the classification accuracy of technical posts on StackOverflow, a technical question-answer forum.

Organization We briefly describe related work in Section 2. Our main contributions are described in Sections 3 and 4. We present an evaluation of our knowledge-base in Section 5. We conclude and identify avenues for future work in Section 6.

[1]Note that, named entities, such as `heap sort`, even though an algorithm are still included in our KB.

2 Related Work

Knowledge-bases have many applications – see for example, Google's Knowledge Graph [18] which is used to provide users with concise summaries to queries about entities, semantic search engines [21], question answering systems such as IBM Watson [7], etc. Our aim is to build a domain-specific knowledge-base (in this case, computer science) which helps in these kinds of applications, but in addition can serve as a resource for learning, including, for example, generating a reading order of concepts in order to learn a new one, automatically generating test or quiz questions about concepts to test student understanding, etc.

Recently, systems which facilitate knowledge-base construction from heterogeneous sources have been proposed. For example, DeepDive [5] aims to consume a large number of heterogeneous data sources for extraction and combines evidence from different sources to output a probabilistic KB. Similarly, Google's Knowledge Vault [6] also aims to fuse data from multiple resources on the Web to construct a KB. Our effort is similar in that we make use of heterogeneous data sources and customise our extractions. However, since our focus is quite narrow and we use very few sources, we do not perform any inferencing.

Entity extraction One of the important aspects of building domain-specific knowledge-bases is that a dictionary of terms that are relevant to the domain should be acquired. It is possible that such dictionaries are already available (for example, lists of people), but for others, we need techniques to build this dictionary. [16] gives an overview of supervised and unsupervised methods to recognize entities from text. We follow a more straightforward approach – we specifically target technology websites and write wrappers to extract a list of entities related to computer science.

Information Extraction Research in information extraction to build knowledge-bases make use of a variety of techniques (see [20] for an overview). In general, information extraction can be done from mostly structured resources such as Wikipedia (see, for example, YAGO [19]) or from unstructured sources (for example, OpenIE [2]) where the relations are not known ahead of time. Moreover, there are rule-based systems such as SystemT [12], using surface patterns and supervised techniques for known relations, distant supervision, etc. (see, for example, Hearst patterns [9] and [15], [4]). We use a mix of these approaches – we formulate different ways to exploit the structured information sources in Wikipedia, and use surface patterns to extract relationships from unstructured sources, such as online books. Some of these techniques provided us with high quality triples, while others failed. We analyze both our successes as well as failures in the paper.

3 TeKnowBase: Acquiring a list of concepts

Our strategy to construct TeKnowbase was to first construct a dictionary of technical concepts and subsequently, to use this dictionary to annotate text and acquire triples. We found that various popular named entity taggers were inadequate in our domain. Therefore, we decided to build our own dictionary from various resources. We converged on three kinds of resources – the oft-used (semi-)structured resource Wikipedia, technology-specific websites Webopedia[2] and TechTarget[3], and subject-specific online textbooks.

Our idea was to use well-structured information from these resources to construct the dictionary of entities. Subsequently, we studied the various structured pieces of information in Wikipedia to extract triples by appropriately engineering our code, and used well-known IE techniques to extract triples from unstructured texts. We describe our ideas in the rest of this section.

[2] `http://www.webopedia.com`
[3] `http://www.techtarget.com/`

3.1 Building a dictionary of technical concepts

There are two main challenges in building a list of technical concepts: i) finding the right resources to extract the concepts (which we previously mentioned), and, ii) reconciling different forms of the same entities (also called the entity resolution problem).

We made use of Wikipedia, technical websites, and online textbooks to build our dictionary. We used Wikipedia's article titles as well as it's category system as a source of concepts. Our corpus of Wikipedia articles consists of all articles under the super category `Computing`. In all, there were approximately 60,000 articles. The titles of each article was considered an entity. Examples entities we found were `Heap_Sort`, `Naive-Bayes_Classifier`, etc.

Our second set of resources were two websites Webopedia and TechTarget. Each website consists of a number technical terms and their definitions in a specific format. From both these websites, we extracted approximately 26,500 entities.

Finally, we extracted entities from the indexes of 6 online textbooks (indexes are also well-structured). These textbooks were specific to the IR and ML domains. The idea can be extended to multiple such online textbooks that are freely available. In all, we extracted approximately 16,500 entities from these textbooks.

While Wikipedia has articles on a number of technical concepts, it is not exhaustive. For example, the terms `average_page_depth` (related to Web Analytics) and `Herbrand_universe` (related to logic) could not be found in Wikipedia, but were found on the technical websites and textbooks respectively.

3.2 Resolving entities

Clearly, the lists of raw concepts extracted from each source have overlaps. We used edit distance to identify and remove duplicates. However, we found that edit distance by itself was not sufficient to resolve all entities, because there are numerous acronyms and abbreviations that are commonly used. Since we wanted to retain both the acronym as well as its expansion as separate entities, we treated the problem of finding (acronym, expansion) pairs as a triple extraction problem – specifically triples for the `synonymOf` relation. Since this involves triple extraction from unstructured sources, we defer the description of our technique to Section 4.

In all, our entity list consists of 85,000 entities after entity resolution.

4 Acquiring relationships between concepts

We classified our relationship extraction task into four types, based on the sources (structured and unstructured) and our prior knowledge of the relations (known relations and unknown relations). By studying our sources and our own idea of the kinds of relations that we wanted in our KB, we made a list of relations to extract (this was our set of known relations). Most importantly, we wanted to construct a taxonomy of concepts, and therefore, the `typeOf` relation was included. It was obvious that we would be unable to list all possible relations, and therefore, we made use of several techniques to acquire unknown relations. We give a brief overview below and describe our techniques in detail from Section 4.1.

- **Structured source, known relations**: As mentioned previously, we were able to recognize, by manual inspection, that the source already contains some of the relationships we want to extract, but organized in a different way. We need to convert one kind of structure (available from the source), into another kind of structure (RDF triples). Wikipedia contains such structured information which can be exploited – specifically, certain types of pages and the organization of the content in a page.
- **Structured source, unknown relations**: Wikipedia also provides us with other kinds of structured information such as template information – that is, tabular data with row headers. These headers potentially describe relationships, but it is not possible to know ahead of time exactly what these relationships are.
- **Unstructured source, known relations**: Unstructured sources are textual sources, such as online textbooks. Our list of entities, constructed as described in Section 3, are of importance here. We could annotate entities in the textual sources and be confident of the correctness of these entities. Subsequently, we find relationships between these entities. We tried the simple technique of using surface patterns for our relations. As we report later in the paper, we were only partially successful in our efforts.

- **Unstructured source, unknown relations**: Finally, we annotated the entities in the textual sources, and ran the open information extraction systems OLLIE [13] to extract *any* relationship between the entities.

4.1 Structured source, known relation

Since our aim was to construct a technical knowledge-base, we manually made a list of relations that our knowledge-base should contain – this is our list of known relations. The relations included the taxonomic relation `typeOf` (as in, ⟨`JPEG typeOf file_format`⟩) and other interesting relationships such as `algorithmFor`, `subTopicOf`, `applicationOf`, `techniqueFor`, etc. In all, we identified 18 relationships that we felt were interesting and formulated techniques to extract them from Wikipedia.

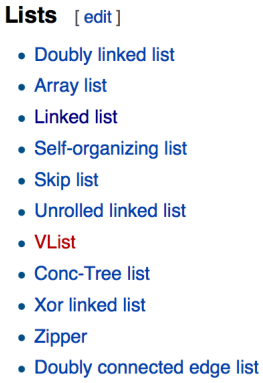

Figure 1: Snippet from "List of Data Structures" page. Extraction of `typeOf` relations from the list structure.

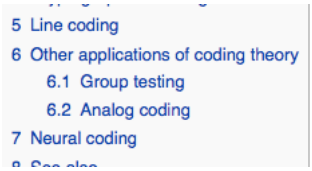

Figure 2: Snippet of the TOC in "Coding theory" page. Extraction of `applicationOf` relation from the list/sublist structure.

Overview pages. We made use of two kinds of structured pages – "List" pages and "Outline" pages (for example, the pages, `List of machine learning concepts`, `Outline of Cryptography`, etc.). These pages organize lists of entities with headings and sub-headings. Extracting this information gives us the relations `typeOf` and `subTopicOf` with good accuracy (see Section 5 for an evaluation of these relations). Figure 1 shows an example of a list page for data structures. We see a list of terms under a heading and can extract triples of the form ⟨`xor_linked_list typeof list`⟩. Further we were able to extract *taxonomic hierarchies* of two levels by relating the headings to the article title. Continuing the previous example, ⟨`List typeOf Data_Structure`⟩ was extracted based on the article title.

Articles on specific topics. These pages refer to discussion on specific topics such as, say, 'Coding Theory'. These pages consist of many structured pieces of information. We made use of three of them of them as follows:

- **The table of contents (TOC)**: From our list of known relations, we searched for keywords within the TOC. If the keyword occurred in an item of the TOC, then the sub-items were likely to be related to it. For example, in Figure 2, the Coding Theory page consists of the following item in its TOC: 'Other applications of coding theory' and this in turn consists of two sub-items 'Group testing' and 'Analog coding'. Since one of the keywords from our known relations is 'application', and the page under consideration is Coding Theory, we extract the triples ⟨`Group_testing applicationOf Coding_theory`⟩ and ⟨`Analog_coding applicationOf Coding_theory`⟩.
- **Section-List within articles**: Next, there are several sub-headings in articles which consist of links to other topics. For example, the page on 'Automated Theorem Proving' consists of a subheading 'Popular Techniques' – this section simply consists of a list of techniques which are linked to their wikipedia page. Since 'technique' is a keyword from our list of known relations, we identify this section-list pattern and acquire triples such as ⟨`Model_checking techniqueFor Automated_Theorem_Proving`⟩.

Source	# found	# triples
List, Outline pages	503	35835
TOC	1838	7412
Section-List	1909	10191
List hierarchies	113	12679
Templates	1139	30434

Table 2: Wikipedia sources from which triples were extracted

- **List hierarchies in articles**: As in the case of "List" pages and "Outline" pages, we make use of list hierarchies in articles to extract the `typeOf` relationships.

Table 2 gives a summary of the number of triples we extracted from each of these sources.

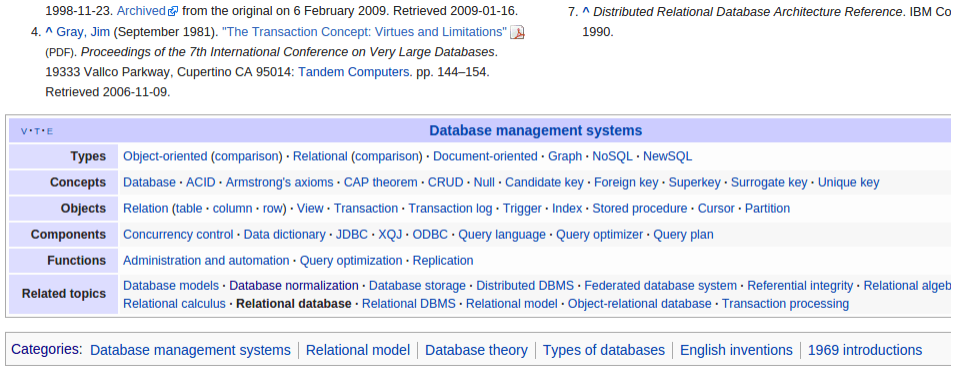
1998-11-23. Archived from the original on 6 February 2009. Retrieved 2009-01-16.
4. ^ Gray, Jim (September 1981). "The Transaction Concept: Virtues and Limitations" (PDF). Proceedings of the 7th International Conference on Very Large Databases. 19333 Vallco Parkway, Cupertino CA 95014: Tandem Computers. pp. 144–154. Retrieved 2006-11-09.
7. ^ Distributed Relational Database Architecture Reference. IBM Co 1990.
Database management systems
Types
Object-oriented (comparison) · Relational (comparison) · Document-oriented · Graph · NoSQL · NewSQL
Concepts
Database · ACID · Armstrong's axioms · CAP theorem · CRUD · Null · Candidate key · Foreign key · Superkey · Surrogate key · Unique key
Objects
Relation (table · column · row) · View · Transaction · Transaction log · Trigger · Index · Stored procedure · Cursor · Partition
Components
Concurrency control · Data dictionary · JDBC · XQJ · ODBC · Query language · Query optimizer · Query plan
Functions
Administration and automation · Query optimization · Replication
Related topics
Database models · Database normalization · Database storage · Distributed DBMS · Federated database system · Referential integrity · Relational algeb
Relational calculus · Relational database · Relational DBMS · Relational model · Object-relational database · Transaction processing
Categories: Database management systems | Relational model | Database theory | Types of databases | English inventions | 1969 introductions

Figure 3: DBMS Template

4.2 Structured source, unknown relations

Continuing with Wikipedia as our nearly structured source, we next set about extracting relationships that were not in our list. For this purpose, we targeted template information that are available in many articles. For example, Figure 3 shows information about Database Management Systems. We made use of the row headers as new relations. For example, in the figure, we have row headings like 'Concepts', 'Objects', 'Functions', etc. which can serve as new relations `conceptOf`, 'functionOf', etc, leading to triples such as ⟨`Query_Optimization functionOf Database_Management_Systems`⟩.

4.3 Unstructured sources

Our unstructured sources include textual description of terms in both Webopedia and Techtarget as well 6 online textbooks related to IR and ML. As previously mentioned, we first annotated the text from all these sources with entities from our entity dictionary. We then tried to extract relationships from them as follows:

Known relations. We formulated simple textual patterns for each of our known relations. For example, for the relation `typeOf`, we used the pattern "is a type of" and for the relation `algorithmFor`, we used the pattern "is an algorithm for". Our search for new triples simply found these patterns in the text and if there were annotated entities around the pattern, then these entities were taken as arguments for the relation. We were successful in identifying the `synonymOf` relation[4] by using the patterns "is abbreviation for", "X (Y)"

[4]Recall from Section 3.2 that we had both abbreviations as well as expansions in our dictionary and wanted to resolve them through identifying the `synonymOf` relation between them.

and "is short for" and extracted a little more than 1,000 such triples. Examples include terms like `JPEG`, which resolved to `Joint Photographic Experts Group`

However, similar to the common experience with other types of corpora, surface patterns alone were not sufficient to extract accurate triples in technical corpora either. Since these extractions were not of high quality, we discarded them from the KB. Improving extractions are part of our future work.

Unknown relations. As our last method of extracting relationships, we made use of the open information extraction tool, OLLIE [13]. Given our annotated textual corpus, we ran OLLIE to find any kind of relationship between entities. We were not very successful in extracting crisp relationships and we need to further study if and how open IE techniques can help us.

Overall, TeKnowbase consists of nearly 100,000 triples. Some basic statistics are shown in Table 3. Examples of relation types and statistics of a few selected relations are given in Table 1

No. of unique entities	85162
No. of unique relations	1326
Most frequent relation	`typeOf` with 44,221
Total no. of triples	98,464
No. of triples extracted from Wikipedia	97,323
No. of triples extracted from unstructured sources	1141

Table 3: TeKnowbase Statistics

5 Evaluation

We performed two kinds of evaluations on TeKnowbase. First, we performed a direct evaluation on the quality of the extractions. We sampled a subset of triples and performed a user evaluation on the accuracy of these triples. Second, we used TeKnowbase in a classification experiment to see if features from this KB could improve the classification accuracy in the style of [8]. We describe each of these experiments below and report our results.

5.1 Experiment 1: Evaluation of Quality

Setup. We chose the top-5 most frequent relations for evaluation. These were: `typeOf`, `terminology`, `synonymOf` `subTopicOf` and `applicationOf`. Together, these five relations constitute about 84% of the triples in our KB. We used stratified sampling to sample from each type of resource. Overall, we sampled 2% of triples corresponding to each relation.

Metrics. Since there is no ground truth against which to evaluate these triples, we relied on user judgement. We asked graduate students to evaluate the triples with the help of the Wikipedia sources if required. Each triple was evaluated by two evaluators and we marked a triple as correct only if both evaluators agreed.

Results and Analysis

Table 4 shows the accuracy of triples for each area. We computed the Wilson interval at 95% confidence for each relation.

The best and worst. On closer examination of these results, we found that we achieved the best results for the `synonym` relation. These triples consisted of both expansions of abbreviations, such as `ALU` and `Arithmetic Logic Unit` as well as alternate terminology such as `Photoshop` and `Adobe Photoshop`.

The best source of extractions are the Wikipedia list pages[5]. In our list of top-5 relations, only 3 were extracted from Wikipedia list pages – `typeOf`, `subtopicOf` and `synonymOf` – and all of them were nearly

[5]Recall that Wikipedia list pages are those which list concepts and typically have an article title starting from `List_of_`

Relation	# Evaluated triples	Accuracy
typeOf	851	$82.5\% \pm 1.2\%$
terminologyOf	606	$93\% \pm 3.3\%$
synonymOf	70	$98\% \pm 0$
subTopicOf	55	$93.5\% \pm 2.5\%$
applicationOf	40	$89.7\% \pm 3.8\%$

Table 4: Evaluation of a subset of triples in TeKnowbase.

100% accurate.

The major source of errors was the `typeOf` extractions from sources other than the list pages above and accounted for nearly 50% of the errors in our evaluation set. Triples extracted from TOCs and Section-lists accounted for many of these errors. Recall from Section 4.1 that we search for a keyword from our list of relations in the TOC items and associate the sublist of that item with the relation corresponding to that keyword. This heuristic did not always work well to identify the correct relation. For example, one of the errors was made when "Game types" was an item in the TOC of the page "Game Theory". It listed "Symmetric/Asymmetric" as a type of game, but we extracted ⟨`Symmetric/Asymmetric typeOf Game_theory`⟩ which is incorrect.

Taxonomy. We specifically analysed the taxonomy (all triples consisting of the `typeOf` relation) since this is an important subset of any KB as well as the largest subset in TeKnowbase. As previously mentioned, our taxonomy consists of over $44,000$ triples and our evaluation yielded an accuracy of $82.5\% \pm 1.2\%$ at 95% confidence. The top two sources of these triples were: Wikipedia lists, List Hierarchies. Around 2000 distinct classes were identified, including, `file formats` (nearly 800 triples), `programming languages` (nearly 700 triples), etc. The accuracy of the `typeOf` triples is affected by different kinds of incorrect extractions from these two sources, even though they account for a very small percentage of the total errors. For example, the list of computer scientists is organized alphabetically and we failed to identify the correct class. This is because one of our heuristics is to use the section header as the class in the "List_of" pages (refer to Figure 1 in Section 4.1).

Other triples were interesting, but not very useful. For example, we have several triples of type `TCP and UDP ports`. While these port numbers by themselves indeed belong to the class `TCP and UDP ports`, they are still just numbers with no indication of what services they are typically used for.

Finally, our taxonomy consists of a limited hierarchy – for example, `XOR linked list`, `List` and `Linear Data Structure` form a hierarchy of length two – however, there are more opportunities to complete this hierarchy. For example, there are several `audio programming languages`, `C-family programming languages`, etc. which are in turn `programming languages`, but are currently not identified as such by our system. This is an important challenge for future work.

5.2 Experiment 2: Classification

In [8], the authors showed how classification accuracy could be improved by generating features from domain specific, ontological knowledge. For example, a document belonging to the class "databases" may not actually contain the term "database", but simply have terms related to databases. If this relationship is explicitly captured in TeKnowbase, then that is a useful feature to add. We conducted a simple experiment to evaluate the use of TeKnowbase in a classification task.

Setup. StackOverflow[6] is a forum for technical discussions. A page in the website consists of a question asked by a user followed by several answers to that question. The question itself may be tagged by the user with several hashtags. The administrators of the site classify the question into one of several technical categories. Our task is to classify a given question *automatically* into a specific technical category.

[6] `stackoverflow.com`

We downloaded the StackOverflow data dump and chose questions from 3 different categories: "databases", "networking", and, "data-structures" We created a corpus of 1500 questions including the title (500 for each category). The category into which the questions were manually classified by the StackOverflow site were taken as the ground truth.

Features generation. First, we annotated the posts with entities from our dictionary. Entities as a whole were treated as a feature and were not broken up into separate words. We generated the following set of features for training.

- **BOW**: Bag of words model (note that entities remained whole).
- **BOW-TKB**: In addition to the words and entities above, for each *entity*, we generated an additional set of features by looking at the relationships the entities participated in. In particular, we used the following relations: `algorithm`, `definition`, `concept`, `topic`, `approach` and `method`. For example, if the entity `run length encoding` occurred in the post, then we added as a feature `data compression` since we have the triple ⟨`run_length_encoding methodOf data_compression`⟩.

Classification algorithms. We trained both a Naive-bayes classifier as well as SVM with each of the feature sets above.

Results. We performed 5-fold cross validation with each classifier and feature set and report the accuracies in Table 5. Clearly, simply adding new features from TeKnowbase helps in improving the accuracies of the classifiers. This result is encouraging and we expect that optimizing the addition of features (for example, coming up with heuristics to decide which relations to use) will result in further gains.

	BOW	**BOW-TKB**
SVM	87.1%	92%
Naive Bayes	88.4%	89.6%

Table 5: Average classification accuracies. Both classifiers show improvement in accuracies with features generated from our KB.

6 Conclusions and Future Work

In this paper, we described the construction of TeKnowbase, a knowledge-base of technical concepts related to computer science. Our approach consisted of two steps – constructing a dictionary of terms related to computer science and to extract relationships among them. We made use of both structured and unstructured information source to extract relationships. Our experiments showed an accuracy of about 88% on a subset of triples. We further used our KB in a classification task and showed how the features generated using the KB can increase classification accuracy.

There are lot of improvements that can be made to our system, purely to increase coverage. We used simple techniques, such as surface patterns, to extract relationships from textual sources. We can try more complex, supervised techniques to do the same. In order to extract unknown relationships, we are interested in exploring open IE techniques in more detail, particularly in identifying interesting and uninteresting relationships.

References

[1] Ashburner, M., Ball, C.A., Blake, J.A., Botstein, D., Butler, H., Cherry, J.M., Davis, A.P., Dolinski, K., Dwight, S.S., Eppig, J.T., Harris, M.A., Hill, D.P., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J.C., Richardson, J.E., Ringwald, M., Rubin, G.M., Sherlock, G.: Gene Ontology: tool for the unification of biology. Nature Genetics 25(1), 25–29 (2000)

[2] Banko, M., Cafarella, M.J., Soderland, S., Broadhead, M., Etzioni, O.: Open Information Extraction from the Web. IJCAI (2007)

[3] Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Hruschka Jr, E.R., Mitchell, T.M.: Toward an Architecture for Never-Ending Language Learning. AAAI (2010)

[4] Carlson, A., Betteridge, J., Wang, R.C., Hruschka Jr, E.R., Mitchell, T.M.: Coupled semi-supervised learning for information extraction. WSDM (2010)

[5] De Sa, C., Ratner, A., Ré, C., 0001, J.S., Wang, F., Wu, S., Zhang, C.: DeepDive - Declarative Knowledge Base Construction. SIGMOD Record 45(1), 60–67 (2016)

[6] Dong, X., Gabrilovich, E., Heitz, G., Horn, W., Lao, N., Murphy, K., Strohmann, T., Sun, S., Zhang, W.: Knowledge vault - a web-scale approach to probabilistic knowledge fusion. KDD pp. 601–610 (2014)

[7] Ferrucci, D.A., 0001, E.W.B., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J.M., Schlaefer, N., Welty, C.A.: Building Watson - An Overview of the DeepQA Project. AI Magazine (2010)

[8] Gabrilovich, E., Markovitch, S.: Feature Generation for Text Categorization Using World Knowledge. IJCAI (2005)

[9] Hearst, M.A.: Automatic Acquisition of Hyponyms from Large Text Corpora. COLING (1992)

[10] Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S., Bizer, C.: DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia. Semantic Web (2015)

[11] Lenat, D.B.: CYC: a large-scale investment in knowledge infrastructure. Communications of the ACM 38(11), 33–38 (1995)

[12] Li, Y., Reiss, F., Chiticariu, L.: SystemT - A Declarative Information Extraction System. ACL (2011)

[13] Mausam, Schmitz, M., Soderland, S., Bart, R., Etzioni, O.: Open Language Learning for Information Extraction. EMNLP-CoNLL (2012)

[14] Miller, G.A.: WordNet - A Lexical Database for English. Commun. ACM (1995)

[15] Mintz, M., Bills, S., Snow, R., Jurafsky, D.: Distant supervision for relation extraction without labeled data. ACL/IJCNLP (2009)

[16] Ren, X., El-Kishky, A., Wang, C., Han, J.: Automatic Entity Recognition and Typing in Massive Text Corpora. WWW (2016)

[17] Seyler, D., Yahya, M., Berberich, K.: Generating Quiz Questions from Knowledge Graphs. WWW (2015)

[18] Singhal, A.: Introducing the knowledge graph: things, not strings. Official Google Blog (2012)

[19] Suchanek, F.M., Kasneci, G., Weikum, G.: Yago - a core of semantic knowledge. WWW (2007)

[20] Suchanek, F.M., Weikum, G.: Knowledge Bases in the Age of Big Data Analytics. PVLDB (2014)

[21] Tummarello, G., Cyganiak, R., Catasta, M., Danielczyk, S., Delbru, R., Decker, S.: Sig.ma - live views on the web of data. WWW (2010)